AD-A094 253   JOHNS HOPKINS UNIV BALTIMORE MD DIV OF BEHAVIORAL BIOLOGY   F/G 12/1
EVENT TIME-SERIES APPLICATIONS TO THE ANALYSIS OF BEHAVIORAL EV--ETC(U)
JAN 81   R L RAY, M J NELLIS, H H EMURIAN   N00014-80-C-0467

UNCLASSIFIED   TR-1   NL

1 of 1
AD A
094253

END
DATE
FILMED
2-81
DTIC

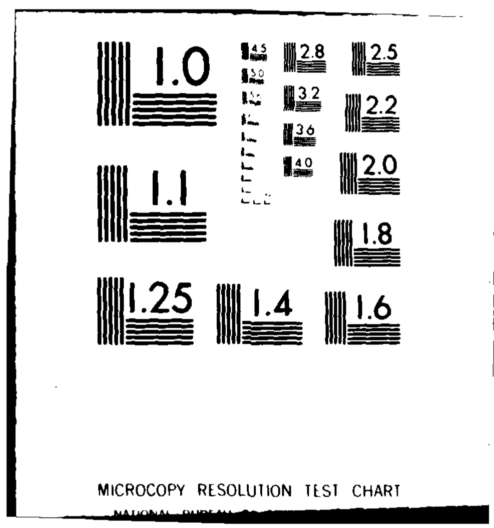
1.0
1.1
1.25
1.4
1.6
1.8
2.0
2.2
2.5
2.8
3.2
3.6
4.0
4.5
5.0
MICROCOPY RESOLUTION TEST CHART

(14) TR-1

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER: TR #1

2. GOVT ACCESSION NO.: AD-A094253

3. RECIPIENT'S CATALOG NUMBER:

4. TITLE (and Subtitle): Event time-series applications to the analysis of behavioral events

5. TYPE OF REPORT & PERIOD COVERED: (9) Technical Report

6. PERFORMING ORG. REPORT NUMBER:

7. AUTHOR(s): (10) R. L. Ray, M. J. Nellis, ~~and~~ H. H. Emurian

8. CONTRACT OR GRANT NUMBER(s): (15) N00014-80-C-0467

9. PERFORMING ORGANIZATION NAME AND ADDRESS:
Division of Behavioral Biology
The Johns Hopkins Univ. Sch. Med.
~~Baltimore, MD~~

10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS: NR 170-910

11. CONTROLLING OFFICE NAME AND ADDRESS:
Organizational Effectiveness Research Programs
Office of Naval Research (Code 452)
Arlington, VA 22217

12. REPORT DATE: (11) 21 Jan[redacted]981

13. NUMBER OF PAGES: 30

14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office): (12) 41

15. SECURITY CLASS. (of this report): UNCLASSIFIED

15a. DECLASSIFICATION DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Event time-series; time-series; quantitative methods.

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

This report discusses the application of event time-series techniques to the analysis of a series of behavioral events. These techniques can be used to supplement more traditional methods of data presentation, such as the cumulative record. The event time-series procedures presented involve the estimation of numerical coefficients and the generation of specialized graphical displays which can be used to characterize the ~

410617

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

Experimental analyses of behavioral phenomena usually emphasize response rate as a datum operationalized by a succession of discrete events in time. Such a series of discrete events in time is called an event time-series. The most popular technique used by behavioral scientists to characterize an event time-series is the cumulative record (Skinner, 1938), although related techniques, such as frequency distributions of inter-event intervals and inter-response-times per opportunity (Anger, 1956), have also been used.

These techniques, particularly the cumulative record, have the advantage that patterns of behavior in individual subjects can be examined and compared in a strictly empirical fashion. One difficulty with the cumulative record, however, is that it often contains so much information that an investigator is burdened to classify subjects or to evaluate the effects of various treatments in a reliable quantitative fashion. Numerical coefficients and specialized graphical displays which express limited aspects of the data can often be useful in overcoming this difficulty, especially when large amounts of data are involved. Accordingly, the purpose of this technical report is to acquaint behavioral scientists with several event time-series techniques which have previously been used by neurophysiologists to study neuronal spike trains and which may be useful in the study of behavioral events.

Renewal Processes

The simplest time-series process of discrete events observed throughout an arbitrary time period is known as a renewal process.

renewal process is one in which the previous history of the process has no effect on its present state of activity, or, in other words, the length of the present inter-event interval is independent of the lengths of any preceding interval (Glaser and Ruchkin, 1976). All of the information necessary to characterize the process is contained in the frequency distribution of between-event intervals.

One familiar renewal process is the Poisson process. In the Poisson process, the probability of an event occurring within any interval of time is proportional to the length of that interval and is independent of the occurrence of previous events generated by the process (Glaser and Ruchkin, 1976). If a Poisson process is operative, the frequency distribution of inter-event intervals is exponential in shape, or in other words, it declines in an exponential and monotonically decreasing fashion from the smallest to largest bin of a frequency distribution histogram. In this case, the single parameter or essential numerical descriptor of the interval data set is the average inter-event interval. This parameter can be computed from the total duration of the data sequence from the first event to the n^{th} event (T_n) by:

$$V = N/T_n \qquad (1)$$

where *V* is the average rate. Glaser and Ruchkin (p. 311) give the details for the construction of a confidence interval for *V*. If an event time-series is generated by a Poisson process, the use of the IRTs/Op. technique with such a series, ignoring sampling fluctuation, will result in a constant IRTs/Op. across class intervals. This results from the fact that

the IRTs/Op. estimate is essentially the ratio of frequencies in successive bins of the frequency distribution histogram, and such a ratio is constant when the frequencies decline exponentially. Exponential frequency distributions have been observed for intervals between lever presses when rats were conditioned to lever press for reinforcements available every three minutes (Mueller, 1950) as well as for inter-spike intervals for spinal interneurones (McGill, 1963).

A more generally useful renewal process is the *gamma process*. This process generates a frequency distribution which can be characterized by two parameters, or two essential numerical quantities, a rate parameter (*V*) and a shape parameter (*r*). The distribution function for the gamma distribution is:

$$P\,(Z) = L\,(r)^{-1} \quad (VZ)^{r-1} \exp\,(Z\,V) \qquad (2)$$

where P (Z) means the probability density of an interval of length Z, *V* is the average rate, *r* is the shape parameter, and L (r) is a mathematical function called the gamma function. The shape parameter determines, as its name implies, the shape of the frequency distribution of intervals. If *r* = 1, the frequency distribution of intervals is exponential, and as *r* approaches ∞, the frequency distribution approaches the normal distribution. Figure 1 shows the shape of the gamma distribution when *r* varies and the scale is held constant. While the gamma distribution is somewhat imposing at first glance, it can be used to model interval frequency distributions with a wide variety of shapes using only two parameters, *r* and *V*. This distribution has been used to describe the frequency

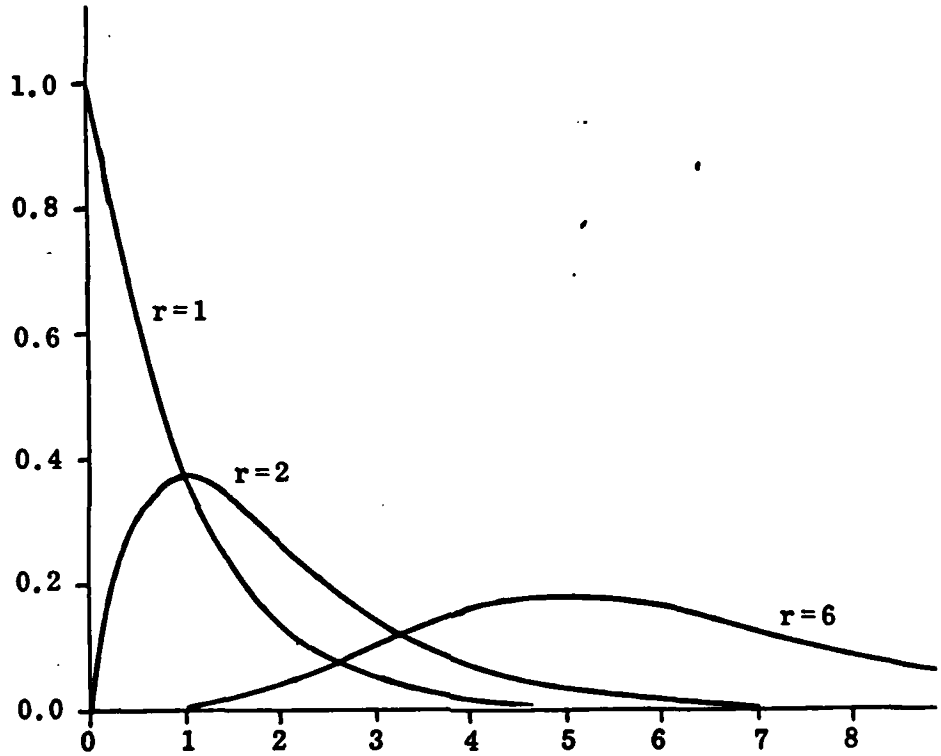

Figure 1. Examples of the gamma distribution for several values of the shape parameter, r. Adapted from Durand (1971).

distribution of inter-spike intervals for certain neurons (Kuffer, Fitzhugh, and Barlow, 1957) and to describe the distribution of runway latencies in rats (McGill, 1963).

The parameters of the gamma distribution can be estimated in several ways (Durand, 1971; Cox and Lewis, 1966; Greenwood and Durand, 1960). One simple way to estimate these parameters is to consult a table provided by Greenwood and Durand (1960). First the arithmetic mean of the intervals (Z) and the mean of the $\log_e$ intervals (ln G) are calculated. The following difference is then obtained:

$$Y = \ln Z - \ln G \tag{3}$$

The obtained value of $\underline{Y}$ is entered in the $\underline{n}$ column of Greenwood and Durand's (1960) Table 1, and the corresponding value of $\underline{np}$ is obtained. When this value is divided by $\underline{Y}$, $\underline{r}$ is obtained and is then divided into $\underline{Z}$ to obtain $\underline{1/V}$. Greenwood and Durand (1960) also give confidence intervals for the estimates of $\underline{r}$ and $\underline{1/V}$. Figure 2 shows the frequency distribution of 320 inter-cigarette intervals collected from a single subject during a ten-day experiment in a residential programmed environment (Nellis, Ray, and Emurian, 1980). The distribution is obviously not exponential. Its mode is not at the smallest frequency bin, and it does not decline in a monotonic fashion. It does, however, resemble a gamma distribution with $\underline{r}>1$. The X's in Figure 2 show the predicted frequencies for each class interval when a gamma distribution is fit to the data.

The gamma distribution is not the only model that may describe an

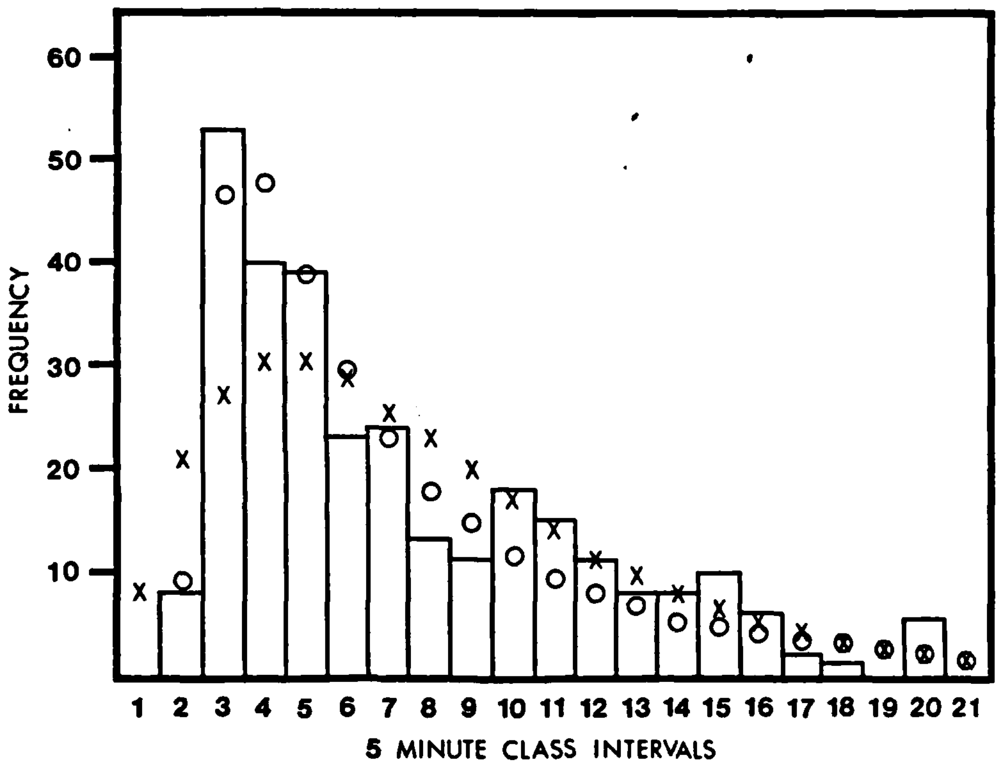

Figure 2. Frequency distribution of inter-cigarette intervals for one subject. The X's correspond to theoretical frequencies predicted by the gamma distribution model and the O's correspond to theoretical frequencies predicted by the lognormal distribution model.

inter-event interval distribution. For example, the inter-event interval distribution may assume the familiar symmetrical, bell-shaped Gaussian or normal form. In this case the two essential parameters which describe the distribution are the mean and standard deviation. In general, however, inter-event interval distributions are not symmetrical. The left-hand tail of the distribution originates at zero or a value greater than zero, and the right-hand tail is strongly skewed. This is the shape seen in Figure 2, and as was shown, the gamma distribution function can be used to model such a distribution.

Another distribution function which can assume a shape comparable with that in Figure 2 is the lognormal distribution. As the name of this distribution suggests, the logarithms of the individual variates of a lognormal distribution are normally or Gaussian distributed. In order to estimate the parameters or essential numerical quantities which describe this distribution, the mean and standard deviation of the logarithmically transformed variables are computed. If the origin of the distribution begins at some value greater than zero, this value is subtracted from each observation prior to logarithmic transformation. If it is desired to compare the fit of the lognormal distribution to an empirical distribution, the following values (Hahn and Shapiro, 1967) are computed:

$$n^* = 1/\hat{S} \tag{4}$$

$$y^* = -X/\hat{S} \tag{5}$$

$$Z^* = -y^* + \ln(X - e) \tag{6}$$

$\underline{X}$ is the value of an observation from the series where $\hat{\underline{S}}$ is the standard deviation of the transformed observations, $\bar{\underline{X}}$ is the mean of the transformed observation, $\underline{e}$ is the origin of the series (≥ 0), and $\underline{Z}$ is a unit normal deviate, or the $\underline{Z}$ found in a standard $\underline{Z}$ table.

The use of (6) allows the cumulative probability associated with a given $\underline{X}$ to be determined, and if the $\underline{X}$'s are the endpoints of the empirical frequency distribution, the observed probabilities and the probabilities predicted by the model can be compared. The open circles in Figure 2 show the predicted frequencies for each class interval. It is apparent that the lognormal distribution fits the empirical distribution of inter-cigarette intervals better than the gamma distribution, at least for this single subject. The parameters of the lognormal distribution were obtained as described above. Each inter-cigarette interval was adjusted by subtracting a constant, the adjusted intervals were logarithmically transformed, and the mean and standard deviation of the adjusted and transformed intervals were computed. The constant subtracted from each inter-cigarette interval was equal to the smallest inter-cigarette interval in the distribution minus one minute, or in this case, seven minutes. This constant was chosen so that the fitted distribution would begin at essentially the smallest value of the empirical distribution. One minute was subtracted from the smallest value of the distribution so that the problem of taking logarithms of zero did not exist.

More sophisticated and efficient procedures exist for fitting lognormal distributions (e.g., Aitchison and Brown, 1957), but the simple

procedure outlined above seemed to produce a fairly good fit to the empirical frequency distribution, certainly a better fit than that obtained with the gamma distribution. While only one subject's data are shown in Figure 2, it seems reasonable to suggest that both gamma and lognormal distributions be considered as potential models for distributions of inter-behavioral events which have shapes similar to that seen in Figure 2.

Serial Dependence

Not all event time-series can be classified as a renewal process. In some event time-series, for example, the length of any particular inter-event interval is dependent on the length of previous inter-event intervals. In most such cases, the type of dependence is assumed to be linear, and it is characterized in terms of the _interval autocorrelation function_. The autocorrelation is defined as follows:

$$r_k = \sum_{1} X_t \, X_{t+k} / \sum_{1} X_t^2 \qquad (4)$$

where n is the total number of intervals, X_t is the deviation of the t^{th} interval from the mean interval and _k_ is the "lag". Lag refers to the number of intervals between the present interval and its comparison interval. For example, when lag=1 autocorrelation is under consideration, the relationship between immediately successive intervals is being examined. When lag=k autocorrelation is under consideration, the linear relationship between successive intervals separated by k-1 intervening intervals is being examined. A plot of r_k for k=1, 2, . . . p is called an autocorrelation function, and p is usually selected to be no greater than

N/10 to N/4 (N=total number of intervals), since fewer and fewer intervals are involved in the computation of r_k as k increases.

When an event time-series is produced by a renewal process, the expected value of all autocorrelations of lag=1 and greater is equal to zero. Because of this, the autocorrelation function is often used to test the hypothesis that the intervals of an event time-series are generated by a renewal process. Because of the extremely non-normal nature of most inter-event interval distributions, Cox and Lewis (1966) recommend that the intervals first be ranked before computing the autocorrelation function. Ranking the intervals also minimizes the effect of gaps in the data on autocorrelation estimates.

Figure 3 shows the ranked autocorrelation estimates for lags 1 to 5 for the inter-cigarette interval series used to construct Figure 2. When n is large, an approximate standard error for R_k is $\sqrt{1/n}$ and $R_k/\sqrt{1/n}$ is approximately distributed as Z, or is unit-normal distributed. None of the five autocorrelations is equal to twice their standard error, and so the hypothesis that their true value is zero cannot be rejected.

One problem with this analysis with the present data set is that the extremely long inter-cigarette intervals corresponding to the subjects' sleep time are included in the series of intervals. It is possible that these long intervals could alter the autocorrelation estimates in such a way as to obscure the "true" pattern of autocorrelation occurring among intervals within each day. For this reason the ranked lag=1 autocorrelation was estimated for each daily series of intervals. The daily

ranked lag=1 autocorrelations ranged from -.363 to +.210 and their weighted mean was -.113. This value was barely significant. The approximate standard error was found by noting that the variance of a weighted combination of independent random variables is equal to the sum of each variable's variance multiplied by the square of the weight applied to that variable. This latter analysis suggests that a small negative relationship exists between the successive inter-cigarette intervals.

If stronger correlation exists between the intervals, a time-series model can sometimes profitably be fit to the series, and the parameters of the model can be used to characterize the behavior of the series. Although the identification and estimation of time-series models are beyond the scope of the present discussion, it is worth noting that time-series can often be modeled with the autoregressive process:

$$X_t = B_1 X_{t-1} + B_2 X_{t-2} + 1 \ldots B_p X_{t-p} + e_t \qquad (5)$$

where the t^{th} interval expressed as a deviation from the mean is $\underline{X_t}$, the $\underline{B}$'s are linear weights, and $\underline{e_t}$ is random error. Weiss, et al. (1966) suggested that interresponse time for rats trained to lever press on a DRL 20 sec schedule could be modeled with an autoregressive process where the value of p was equal to 1. Equation (5) indicates that $\underline{X_t}$ is equal to a weighted combination of $\underline{P}$ previous intervals plus random error. It is possible to estimate the autoregression weights of coefficients with ordinary multiple regression computer programs, and the selection of the number of lagged values of X_t to be used in the autoregression equation can be determined with standard stepwise selection procedures. Other methods

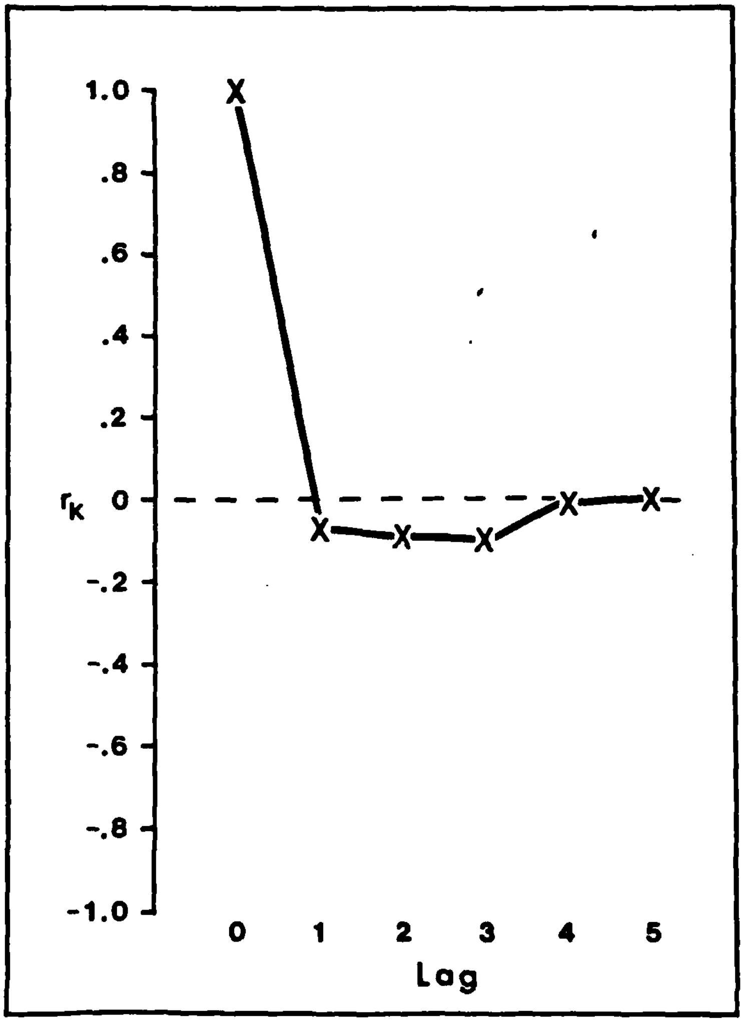

Figure 3. Rank autocorrelations for lags 0 to 5 for the inter-cigarette interval series. Note that the lag-0 autocorrelation is always equal to 1.0.

are more commonly used, however. The prospective user of time-series analysis should consult specialized sources before undertaking this task (*e.g.*, Jones, 1976; Box and Jenkins, 1970; Chatfield, 1976).

The foregoing discussion dealt with the autocorrelation between the intervals separating the successive events of an event time-series. Sometimes the autocorrelation between the successive events themselves is examined. Figure 4 shows how the autocorrelation between events can be estimated. In this figure the horizontal axis is time and each vertical line marks the occurrence in time of a behavioral event. The event autocorrelation at a one-minute lag can be generated by displacing the figure one minute in time and counting the number of times that events from the two figures coincide. This frequency of coincidence is proportional to the event autocorrelation with a one-minute lag. This procedure can be repeated for a wide variety of lags and the plot of the resulting frequency of coincidences has been called the post-event histogram, the event autocorrelogram (Sayers, 1971) and the intensity function (Cox and Lewis, 1966).

The post-event histogram represents the probability of events occurring at various times following an arbitrarily selected event. Another way to think of this technique is as a probability forecasting function. When a behavioral event occurs, the post-event histogram can be used to predict the probability of subsequent events occurring at various times afterwards. Figure 5 shows the post-event histograms for the cigarette event series, and for comparison, the post-event histograms for

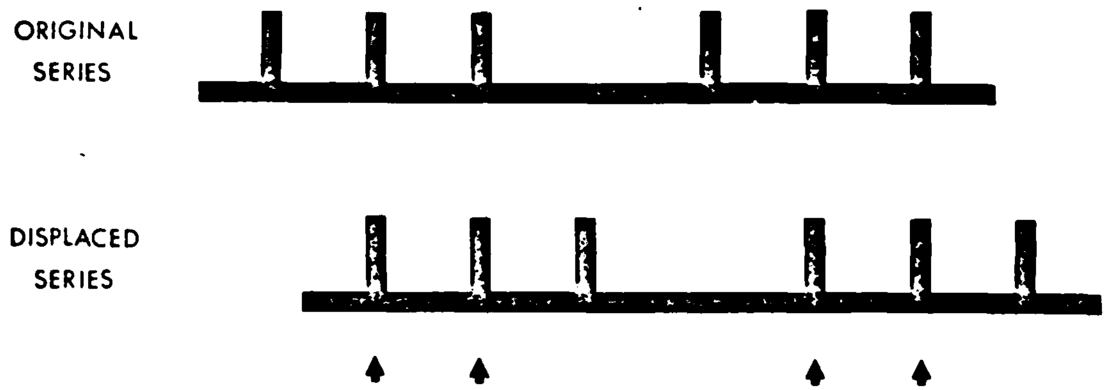

Figure 4. Graph of a hypothetical event-series and of the same series displaced ahead in time by one minute. Arrows mark coincidences between events of the original and displaced series. The coincidences of each lag are summed to produce the post-event histogram.

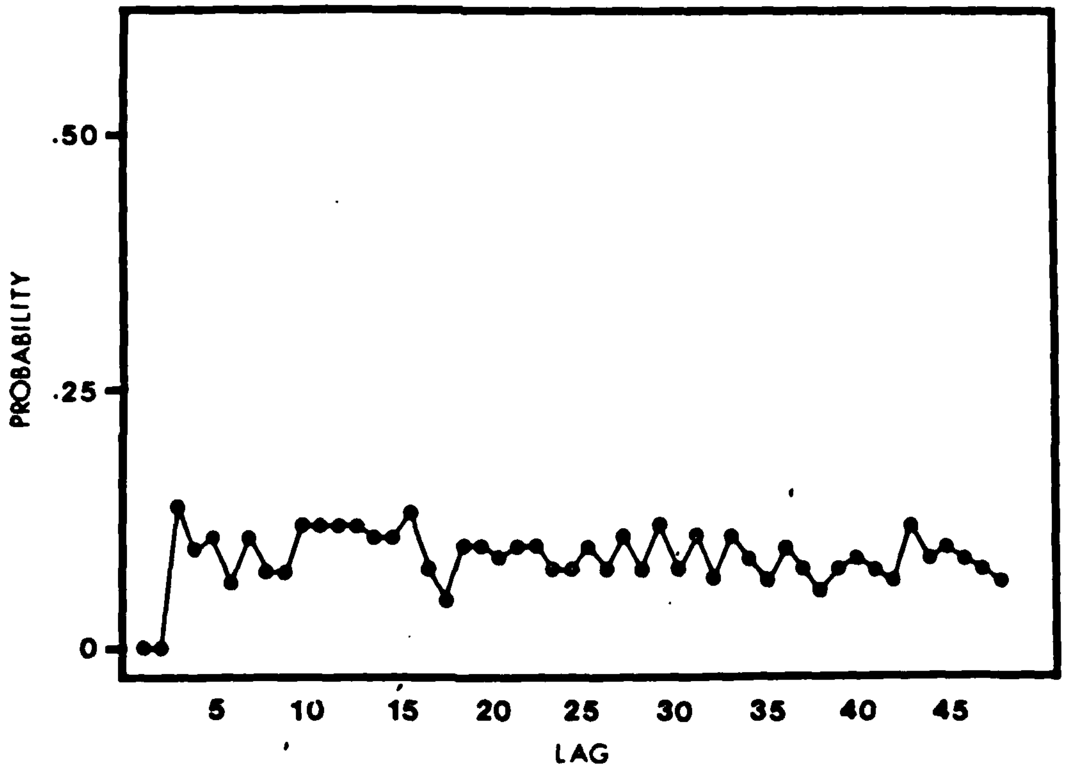

Figure 5a. Post-event histogram (presented as a frequency polygon) for a series of inter-cigarette intervals. The numbers along the abscissa are in units of five minutes.

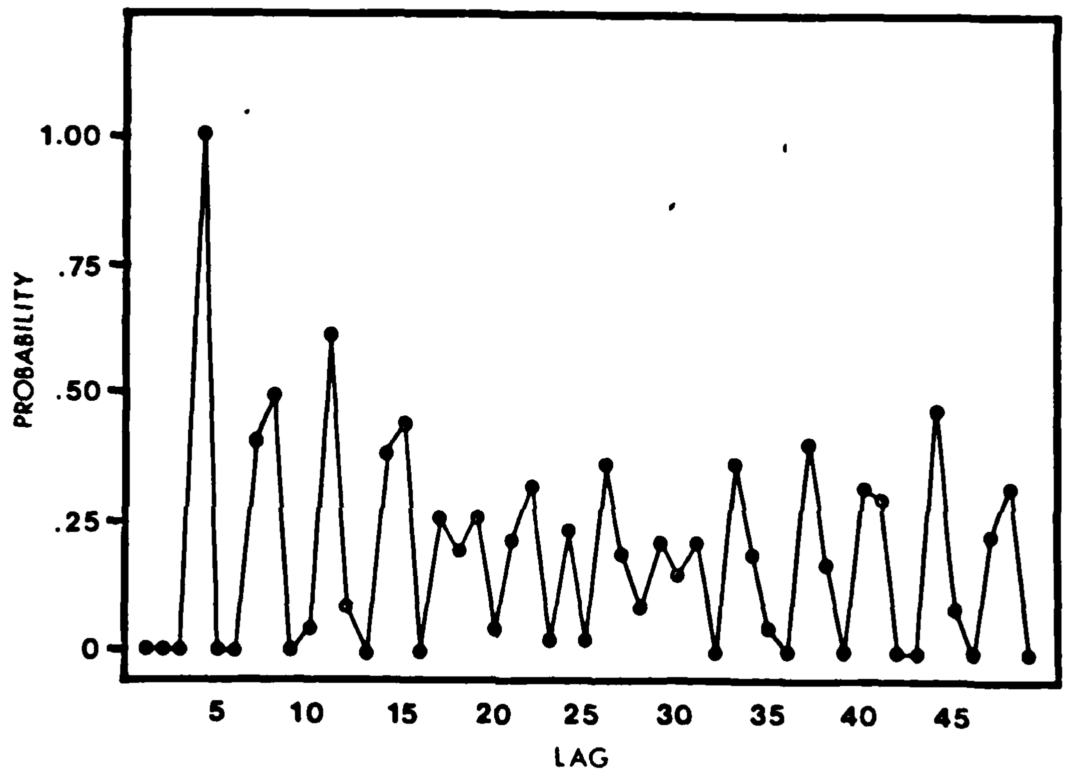

Figure 5b. Post-event histogram (presented as a frequency polygon) for a series of cardiac inter-beat intervals. The numbers along the abscissa are in units of .2 sec.

the cigarette event series, and for comparison, the post-event histogram for a series of sixty successive heart beats. The post-event histogram for cigarette events rises suddenly at 20 minutes and stays more or less constant, with minor irregularities, for the remaining 225 minutes. In contrast the post-event histogram for the heart beats shows a marked periodicity, alternating between probabilities of zero and close to 1.0 for much of the histogram. This periodicity in the post-event histogram is characteristic of regularly occurring processes. It should be noted, however, that periodicities in the post-event histogram are not, in themselves, evidence for autocorrelation among the inter-event intervals. Instead, the periodicity is often due to the shape of the inter-event interval distribution. Distributions with a prominent model interval tend to produce "peaky" post-event histograms.

The actual production of the post-event histograms in Figure 5 was not achieved with the computational procedure described above. Instead the following alternative computational procedure was used. The n(n-1)/2 possible intervals between all events were measured and a frequency distribution was constructed using these intervals. While this is tedious with hand calculations, it is simple to implement with a computer. If the frequency of intervals falling within each bin of the frequency distribution is divided by n (the total number of intervals in the event time-series) the transformed frequencies can be interpreted as probabilities. If the average rate is subtracted from the transformed frequencies, the resulting values can be interpreted as the autocovariances among events at various lags. Inspection of the post-event histogram of

inter-cigarette intervals revealed that the sequence of events is fairly irregular, after an initial latency period, and the probability of subsequent events given the occurrence of any arbitrary event varies from only 10% to 15% over various time lags. It should be noted that if marked periodicity is seen in a post-event histogram, spectral analysis can be used to find the dominant frequency of the periodicity and to determine how much of the variance in the event series is accounted for by this frequency (see Glaser and Ruchkin, 1976; Sayers, 1971; Cox and Lewis, 1961).

Relationships Between Event Time-Series

The relationship between two behavioral event time-series has often been investigated by plotting the cumulative records of the several event time-series concurrently (*e.g.*, Catania, 1966). While this technique allows the possibility of discerning a wide variety of possible relationships between the two series, this technique, as in the univariate or single-series case, can produce so much information that statements about the presence and nature of dependencies between several series become very difficult.

When it is desirable to determine if occurrences of behavioral events in one time-series are related to the occurrences of events in another time-series, the cross-interval histogram is often useful (Perkel, Gerstein, and Moore, 1967; Sayers, 1967). Figure 6 shows a diagrammatic representation of two concurrent event time-series. Each vertical bar represents the occurrence of a behavioral event and the horizontal axis represents time. After viewing such a figure, the question naturally

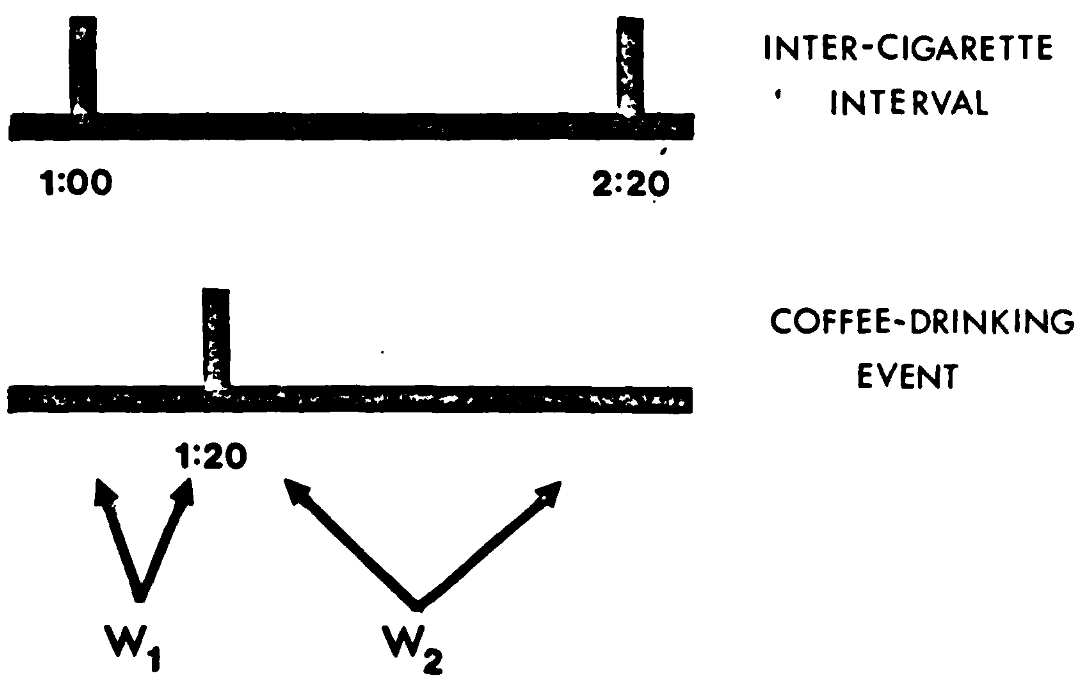

Figure 6. Graph showing an inter-cigarette interval from a cigarette-smoking event-series and a coffee-drinking event falling within the inter-cigarette interval. W_1 corresponds to the backward waiting time and W_2 corresponds to the forward waiting time.

arises as to whether the occurrences in time of the events of the two series are related. The cross-event histogram technique is designed to deal with this question.

The cross-event histogram technique works by focusing on inter-event intervals of one of the two series. After identifying an inter-event interval, the data analyst, or computer, looks to see if an event of the other event time-series fell within the inter-event interval. In Figure 6, for example, an inter-cigarette interval begins at 1:00 and ends at 2:20. A coffee-drinking event occurs within this interval at 1:20. The next step in the analysis is to measure the interval between the beginning of the inter-cigarette interval and the time of the coffee-drinking event (W_1) and the interval between the time of the coffee-drinking event and the end of the inter-cigarette interval (W_2). The interval W_1 is called the backward waiting time and the interval W_2 is called the forward waiting time. Assuming that our two series are independent, <u>there will be no particular place in the inter-cigarette interval that the coffee-drinking event is likely to fall</u>. *In more precise terms, if we measure the* W_1 and W_2 intervals for many inter-cigarette intervals, the frequency distributions of intervals W_1 and W_2 should not differ systematically, given that the two event time-series are independent. Conversely, if the frequency distributions of W_1 and W_2 do appear *systematically different, it is probable* that the two series are not independent of one another. In practice, a great many inter-cigarette intervals would be analyzed. All possible W_1 and W_2 waiting time intervals would be measured and used to construct a histogram like that shown in Figure 7.

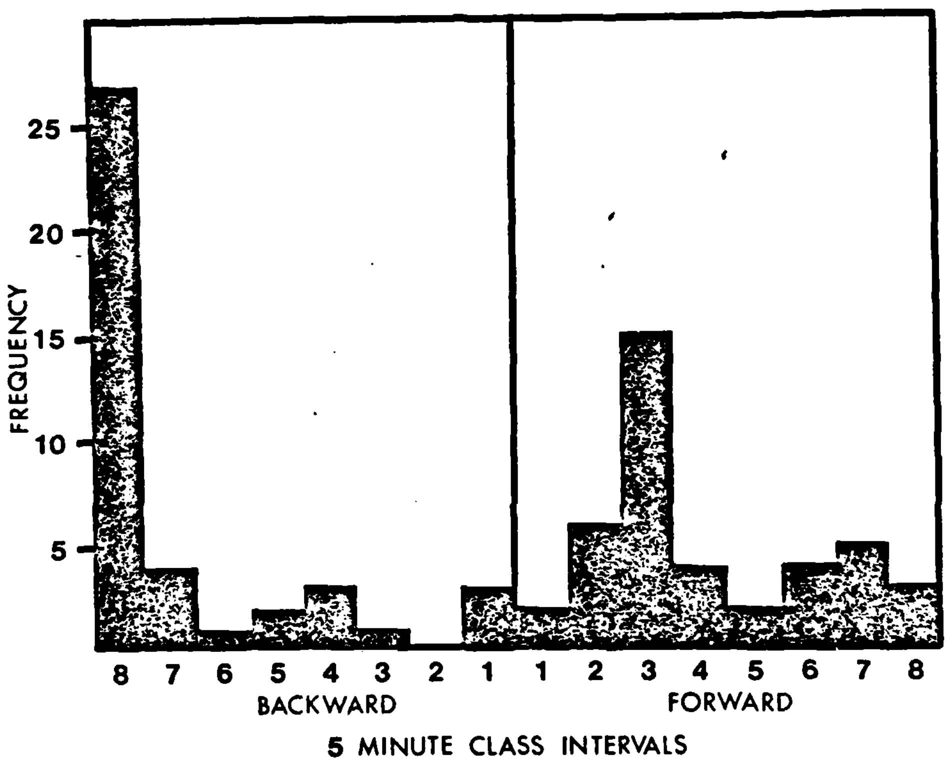

Figure 7. Cross-interval histogram for coffee-drinking events falling within the inter-cigarette intervals.

Since the use of the cross-interval histogram requires arbitrarily choosing one of the two series for examination of the inter-event intervals, it is usually recommended that the procedure be performed twice, once for the inter-event intervals of each series.

The height of each bar on the right-hand side of Figure 7 represents the number of forward waiting times that fall in each class interval, and the height of each bar on the left-hand side of the figure represents the number of backward waiting time intervals that fall in each class interval. When the two event time-series are independent, the left-hand and right-hand sides of the cross-interval histogram should be mirror images of one another. In order to decide whether the two sides of the cross-interval histogram differ by more than chance, Sayer (1967) recommends the use of a significance test which compares two frequency distributions, such as a chi-square test. The outcome of this test must be interpreted with caution, however, since one of its basic assumptions is that the two distributions being compared are independent.

Figure 7 shows the cross-interval histogram for coffee-drinking events falling within inter-cigarette intervals. The distributions of the forward and backward waiting times are clearly different. The backward waiting times tend to cluster around the class interval of $\geq$ 35 minutes, or, in other words, the time between the beginning of the inter-cigarette interval and the coffee-drinking event tends to be long. In contrast, the forward waiting times tend to be short. In other words, the coffee-drinking event tends to occur near the end of the inter-cigarette interval. Figure 8

shows the cross-interval histogram for cigarette-smoking events occurring within inter-coffee intervals. This figure shows that the backward waiting times for cigarettes within inter-coffee intervals tend to be shorter than the forward waiting times.

The cross-interval histogram technique only indicates whether it is likely that two event time-series are independent. If they do appear to be independent, the process responsible for generating the dependency must be investigated with other techniques. One simple way in which dependencies can be generated between two event time-series is the case when the two series are "in phase". Two event time-series are in phase when the backward waiting time of one series is a constant fraction of the inter-event interval of the other series. Figure 9 shows the distribution of the phase fractions across all subjects with respect to coffee-drinking events falling within inter-cigarette intervals. This figure shows that 75% of the phase fractions are greater than 50%. In other words, the coffee-drinking event tends to occur in the last half of the interval. This, of course, would be expected from examination of Figure 7. The phase fractions are not constant, however, but show a great deal of variability.

Figure 10 shows the distribution of phase fractions for cigarette-smoking events falling within inter-coffee intervals. This figure shows the converse of Figure 7. Phase-fractions for cigarette-smoking events falling within inter-coffee intervals tend to be small, the greatest number falling within 0% to 25%. The cross-interval histogram and phase fraction histogram for cigarette events within intercoffee intervals show less

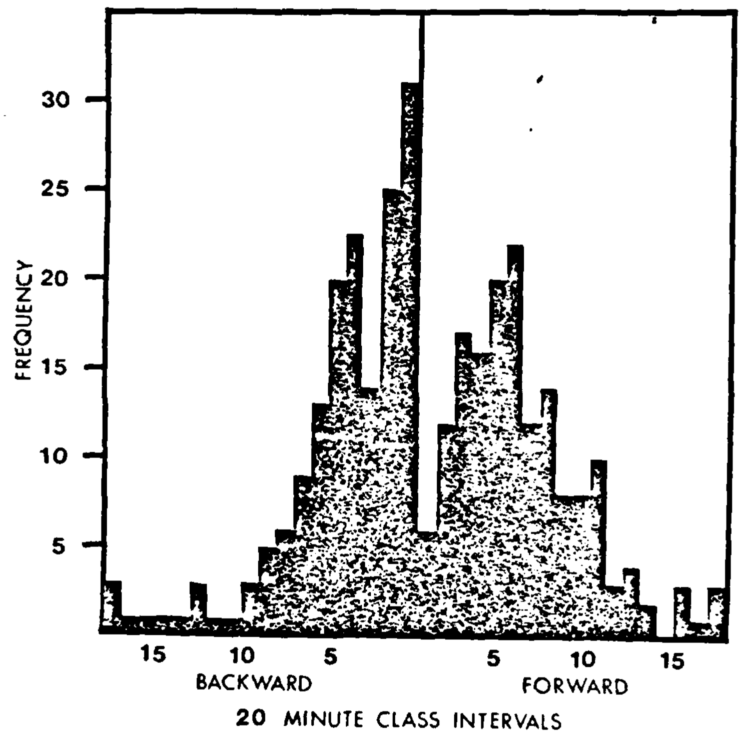

Figure 8. Cross-interval histogram for cigarette-smoking events falling within the inter-coffee intervals.

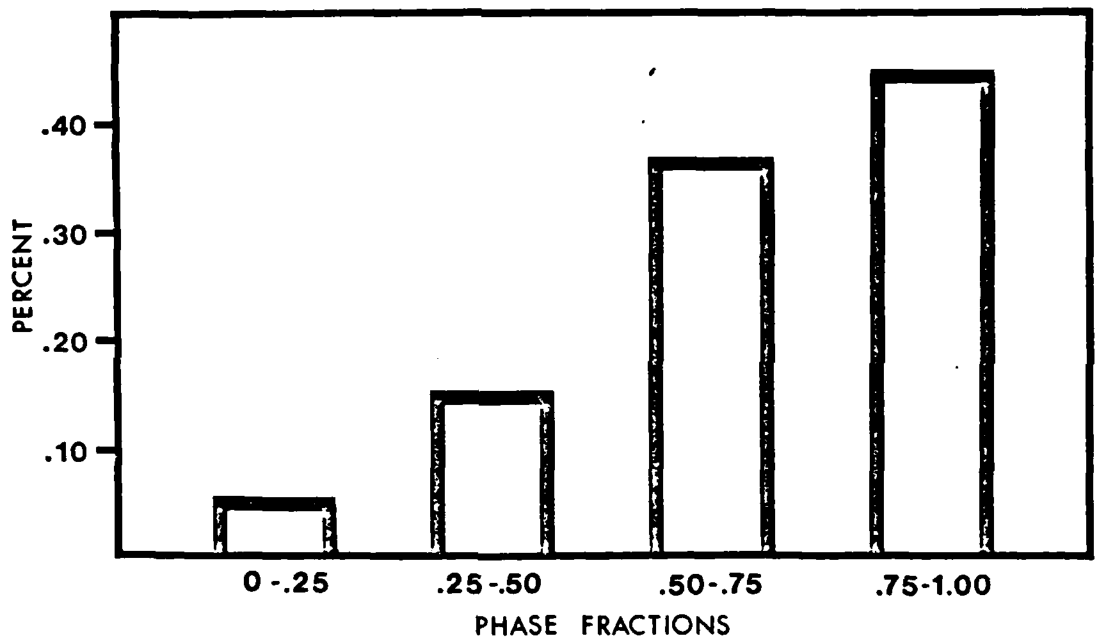

Figure 9. Phase fraction histogram for coffee-drinking events falling within the inter-cigarette intervals.

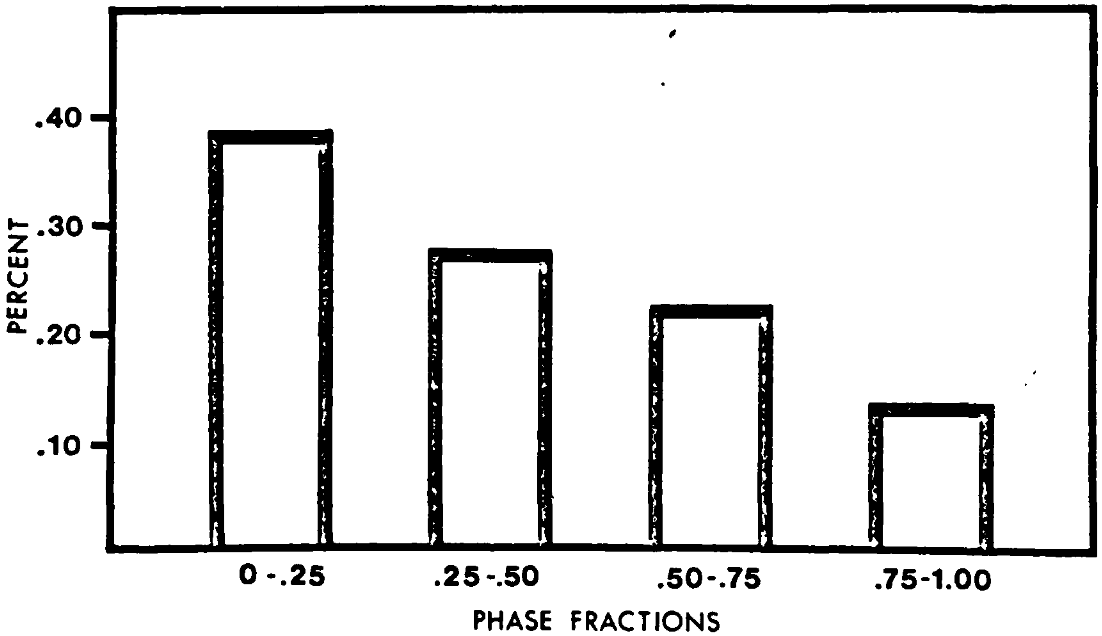

Figure 10. Phase fraction histogram for cigarette-smoking events falling within the inter-coffee intervals.

evidence of dependencies than the corresponding histograms for coffee-drinking events falling within inter-cigarette intervals. This difference is due to the fact that the rate of cigarette smoking is much higher than the rate of coffee drinking. Many cigarettes occur within each inter-coffee interval, forcing a more even distribution of waiting times and phase fractions than that seen for the corresponding histograms involving coffee-drinking events within inter-cigarette intervals.

If the two event time-series are independent of one another, the frequency distributions of phase-fractions, disregarding sampling fluctuation, should be flat, or in other words, show a uniform distribution. This suggests that a chi-square goodness of fit test could be used to test the hypothesis of dependence between the two series by comparing the obtained distribution with a theoretic distribution having equal frequencies at each class interval. Again, as in the case of the cross-interval histogram, the results of this test must be interpreted with caution since its assumptions may not be satisfied with the type of data under consideration. There is nothing to prevent the investigator, however, from computing a chi-square for either type of histogram for the purpose of comparison and classification of subjects. In the present case larger chi-square values indicate a greater degree of dependency between the two series than small values.

The cross-interval histogram and the phase-fraction histogram are useful for detecting simple dependencies between two event time-series. An easily computed chi-square value can be determined which characterizes the

degree to which the shape of these histograms deviates from that which would result if the two series were independent. These are not the only techniques available for characterizing and detecting the dependencies between two event time-series, however. Although these other more complex techniques are beyond the scope of this discussion, they are described in Glaser and Ruchkin (1976) and Sayers (1971).

Discussion

When a series of behavioral events is observed, the dynamic properties of that series and its possible relationship with other series are often of interest. While the most successful procedure used to deal with problems of this nature has been the cumulative record, techniques previously used to analyze series of neuronal spikes, event time-series procedures, can sometimes be used to reveal aspects of the data which are difficult to see in a cumulative record. These techniques involve determining, for a single series, the nature of the distribution of intervals in terms of both the shape of the distribution and presence or absence of sequential dependencies among intervals or events. When possible relationships between several series are at question, the distribution of intervals across series, forward and backward waiting times, and the ratio of backward waiting time to inter-event interval (the phase fraction) can be examined.

The techniques described in this paper are best applied to fairly long series of events and to series which are "steady-state" in nature, or which do not exhibit any noticeable trend or change in variance. The presence of

trends in an event time-series distorts the shape of the distribution of intervals. The trend creates serial correlation among successive intervals which obscures the nature of the underlying process.

The examples cited within this technical report were chosen to illustrate the event time-series techniques with comparatively simple series of events and straightforward response sequences which are highly visable and easily assessed. The techniques themselves, however, have obvious applicability to the analysis of series and interdependencies which may characterize complex organizational systems such as the steady-state operational performance of a team unit whose mission requires coordinated responding among team members as a function of external task demands and individual member contributions to mission goals. It is anticipated that a more comprehensive appreciation of these methodological applications will occur in terms of their relevance to the objectives of the current research program with particular reference to investigations of team performance effectiveness.

References

Aitchison, J., & Brown, J.A.C. The Lognormal Distribution. Cambridge: Cambridge University Press, 1957.

Anger, D. A. The dependence of interresponse times upon the relative reinforcement of different interresponse times. Journal of Experimental Psychology, 1956, 52, 145-161.

Box, G.E.P., & Jenkins, G.M. Time-Series Analysis: Forecasting and Control. San Francisco: Holden-Day, 1970.

Catania, A.C. Concurrent operants. In W.K. Honig (Ed.), Operant Behavior: Areas of Research and Application. New York: Appleton-Century-Crofts, 1966, 213-270.

Chatfield, C. The Analysis of Time-Series: Theory and Practice. London: Chapman and Hall, 1975.

Cox, D.R., & Lewis, P.A.W. The Statistical Analysis of a Series of Events. London: Methuan, 1966.

Durand, D. Stable Chaos: An Introduction to Statistical Control. Morristown: General Learning Corporation, 1971.

Glaser, E.M., & Ruchkin, D.S. Principles of Neurobiological Signal Analysis. New York: Academic Press, 1976.

Greenwood, J.A., & Durand, D. Aids for fitting the gamma distribution by maximum likelihood. Technometrics, 1960, 2, 55-65.

Hahn, G.J., & Shapiro, S.S. Statistical Models in Engineering. New York: John Wiley & Sons, 1967.

Jones, R.H. Fitting autoregressions. Journal of the American Statistical Association, 1975, 70, 590-592.

Kuffler, S.W., Fitzhugh, R., & Barlow, H.B. Maintained activity in the cats retina in light and darkness. Journal of General Physiology, 1957, 49, 683-702.

McGill, W. Stochastic latency distributions. In R.D. Luce, R.R. Bush, & E. Galanter (Eds.), Handbook of Mathematical Psychology. Vol. I. New York: John Wiley, 1966, 309-360.

Mueller, C.G. Theoretical relationships among some measures of conditioning. Proceedings of the National Academy of Sciences, 1950, 36, 123-130.

Nellis, M. J., Ray, R.L., & Emurian, H.E. Event time-series relationship between cigarette smoking and coffee drinking. Presented at the annual meetings of the Eastern Psychological Association, 1981.

Perkel, D., Gerstein, G.L., & Moore, G.P. Neuronal spike trains: I. The single spike train. Biophysical Journal, 1967, 1, 391-418.

Perkel, D., Gerstein, G.L., & Moore, G.P. Neuronal spike trains: II. Simultaneous spike trains. Biophysical Journal, 1967, 1, 419-440.

Sayers, B. Inferring significance from biological signals. In M. Clynes and J. Milsum (Eds.) Biomedical Engineering Systems. New York: McGraw-Hill, 1970, 84-163.

Skinner, B.F. The Behavior of Organisms. New York: Appleton-Century-Crofts, 1938.

Weiss, B., Laties, V., Siegel, L., & Goldstein, D. A computer analysis of serial interactions in spaced responding. Journal of the Experimental Analysis of Behavior, 1966, 9, 619-626.

DISTRIBUTION LIST

Office of Naval Research
Code 452
800 North Quincy Street
Arlington, VA 22217

Director
U.S. Naval Research Laboratory
Washington, D.C. 20390
ATTN: Technical Information Div.

Defense Documentation Center
Building 5
Cameron Station
Alexandria, VA 22314

Psychologist, ONR Branch Office
495 Summer Street
Boston, MA 02210

Commanding Officer
ONR Branch Office
536 South Clark Street
Chicago, IL 60605

Dr. David G. Bowers
Institute for Social Research
University of Michigan
Ann Arbor, Michigan 48106

Dr. J. Richard Hackman
Administrative Sciences
Yale University
56 Hillhouse Avenue
New Haven, Connecticut 06520

Dr. S. B. Sells
Texas Christian University
Fort Worth, TX 76129

Dr. Meredith Crawford
5605 Montgomery Street
Chevy Chase, MD 20015

Science & Technology Division
Library of Congress
Washington, D.C. 20540

Commanding Officer
ONR Branch Office
1030 East Green Street
Pasadena, CA 91106

Psychologist, ONR Branch Office
1030 East Green Street
Pasadena, CA 91106

Research Psychologist, ONR
ONR Branch Office
536 South Clark Street
Chicago, IL 60605

Dr. Earl A. Alluisi
Old Dominion University
Research Foundation
Norfolk, VA 23508

Dr. Samuel L. Gaertner
Department of Psychology
University of Delaware
220 Wolf Hall
Newark, Delaware 19711

Dr. Irwin Sarason
Department of Psychology
University of Washington
Seattle, WA 98195

Dr. Lennart Levi, Director
Lab. Clinical Stress Research, FACK
S-104 01 Stockholm, Sweden

Dr. Edwin Hollander
Department of Psychology
State University of New York
at Buffalo
4230 Ridge Lea Road
Buffalo, NY 14226

-8